

BIJECTIONS AND SYMMETRIES FOR THE FACTORIZATIONS OF THE LONG CYCLE

OLIVIER BERNARDI AND ALEJANDRO H. MORALES

ABSTRACT. We study the factorizations of the permutation $(1, 2, \dots, n)$ into k factors of given cycle types. Using representation theory, Jackson obtained for each k an elegant formula for counting these factorizations according to the number of cycles of each factor. In the cases $k = 2, 3$ Schaeffer and Vassilieva gave a combinatorial proof of Jackson's formula, and Morales and Vassilieva obtained more refined formulas exhibiting a surprising symmetry property. These counting results are indicative of a rich combinatorial theory which has remained elusive to this point, and it is the goal of this article to establish a series of bijections which unveil some of the combinatorial properties of the factorizations of $(1, 2, \dots, n)$ into k factors for all k . We thereby obtain refinements of Jackson's formulas which extend the cases $k = 2, 3$ treated by Morales and Vassilieva. Our bijections are described in terms of "constellations", which are graphs embedded in surfaces encoding the transitive factorizations of permutations.

1. INTRODUCTION

We consider the problem of enumerating the factorizations of the permutation $(1, 2, \dots, n)$ into k factors according to the cycle type of each factor. In [6] Jackson established a remarkable *counting formula* (analogous to the celebrated Harer-Zagier formula [5]) characterizing the generating function of the factorizations of the long cycle according to the number of cycles of each factor. A combinatorial proof was subsequently given for the cases $k = 2, 3$ by Schaeffer and Vassilieva [13, 12]. Building on these bijections, Morales and Vassilieva also established for $k = 2, 3$ a formula for the generating function of factorizations of $(1, 2, \dots, n)$ counted according to the cycle type of each factor [10, 11]. These formulas display a surprising *symmetry property* which has remained unexplained so far.

In this article we explore the combinatorics of the factorizations of the permutation $(1, 2, \dots, n)$ through a series of bijections. All our bijections are described in terms of *maps* and *constellations* which are graphs embedded in surfaces encoding the transitive factorizations in the symmetric group (see Section 2 for definitions). A summary of our bijections is illustrated

Date: December 23, 2011.

O. Bernardi acknowledges support from NSF grant DMS-1068626, ERC ExploreMaps and ANR A3.




FIGURE 1. Summary of the bijections presented in this article. We start with the classical encoding of the factorizations of $(1, 2, \dots, n)$ by cacti (here $k = n = 3$). We then establish a bijection Φ between vertex-colored cacti and tree-rooted constellations (Section 3). We then characterize the dual of tree-rooted constellations and obtain a correspondence with nebulas (Section 5). Lastly we establish a bijection Ψ between nebulas and valid biddings.

in Figure 1. Our first bijection gives an encoding of the factorizations of the permutation $(1, 2, \dots, n)$ into *tree-rooted k -constellations*. This encoding allows one to easily prove the case $k = 2$ of Jackson's counting formula, as well as to establish the symmetry property for all $k \geq 2$. However for $k \geq 3$, the tree-rooted k -constellations are still uneasy to count and we give further bijections. Eventually, we show bijectively that proving Jackson's counting formula reduces to proving an intriguing probabilistic statement (see Theorem 1.6). In Section 7 we prove this probabilistic statement in the cases $k = 2, 3, 4$ (thereby proving Jackson's counting formula for these cases) but the cases $k > 4$ shall be treated (along with similar probabilistic statements) in a separate paper [3]. Before describing our results further we need to review the literature.

Enumerative results about the factorizations of the long cycle.

Given k partitions $\lambda^{(1)}, \dots, \lambda^{(k)}$ of n , it is a classical problem to determine the number $\kappa(\lambda^{(1)}, \dots, \lambda^{(k)})$ of factorizations $\pi_1 \circ \pi_2 \circ \dots \circ \pi_k = (1, 2, \dots, n)$ such that the permutation π_t has cycle type $\lambda^{(t)}$ for all $t \in \{1, \dots, k\}$. By the general theory of group representations, the *connection coefficients* $\kappa(\lambda^{(1)}, \dots, \lambda^{(k)})$ can be expressed in terms of the characters of the symmetric group, but this expression is not really explicit even for $k = 2$. However, Jackson established in [6] a remarkable formula for the generating function of factorizations counted according to the number of cycles of the factors, namely,

$$(1.1) \quad \sum_{\pi_1 \circ \dots \circ \pi_k = (1, 2, \dots, n)} \prod_{i=1}^k x_i^{\ell(\pi_i)} = \sum_{1 \leq p_1, \dots, p_k \leq n} \prod_{i=1}^k \binom{x_i}{p_i} n^{k-1} M_{p_1-1, \dots, p_k-1}^{n-1}$$

where $\ell(\pi)$ is the number of cycles of the permutation π , and M_{p_1, \dots, p_k}^n is the coefficient of $x_1^{p_1} \dots x_k^{p_k}$ in the polynomial $(\prod_{i=1}^k (1 + x_i) - \prod_{i=1}^k x_i)^n$.

Jackson's formula can equivalently be stated in terms of *colored factorizations*.

Definition 1.1. *Given positives integers p_1, \dots, p_k , a (p_1, \dots, p_k) -colored factorization of $(1, 2, \dots, n)$ is a tuple $(\pi_1, \dots, \pi_k, \phi_1, \dots, \phi_k)$, where π_1, \dots, π_k are permutations of $[n] := \{1, \dots, n\}$ such that $\pi_1 \circ \dots \circ \pi_k = (1, 2, \dots, n)$ and for all $t \in [k]$, ϕ_t is a surjective mapping from $[n]$ to $[p_t]$ such that $\phi_t(a) = \phi_t(b)$ if a, b are in the same cycle of π_t . In other words, the mapping ϕ_t can be seen as a coloring of the cycles of the permutation π_t with colors in $[p_t]$ and we want all the colors to be used.*

It is easy to see that (1.1) is equivalent to the following theorem.

Theorem 1.2 (Jackson's counting formula [6]). *The number C_{p_1, \dots, p_k}^n of (p_1, \dots, p_k) -colored factorizations of the permutation $(1, 2, \dots, n)$ is equal to*

$$(1.2) \quad n!^{k-1} M_{p_1-1, \dots, p_k-1}^{n-1},$$

where $M_{p_1, \dots, p_k}^n = [x_1^{p_1} \dots x_k^{p_k}] (\prod_{i=1}^k (1 + x_i) - \prod_{i=1}^k x_i)^n$ is the cardinality of the set $\mathcal{M}_{p_1, \dots, p_k}^n$ of n -tuples (R_1, \dots, R_n) of strict subsets R_t of $[k]$ such that each integer $t \in [k]$ appears in exactly p_t of the subsets R_1, \dots, R_n .

The original proof of Theorem 1.2 in [6] is based on the representation theory of the symmetric group. Bijections explaining the cases $k = 2, 3$ were subsequently given by Schaeffer and Vassilieva [13, 12]. The case $k = 2$ of Theorem 1.2 is actually closely related to the celebrated Harer-Zagier formula [5], which was proved bijectively by Goulden and Nica [4]. In Section 3 we shall give a bijection which extends the results in [13, 12] to arbitrary k (however, for a general k , this bijection does not directly imply Theorem 1.2).

We now consider a refined enumeration problem. Let $\gamma^{(1)}, \dots, \gamma^{(k)}$ be compositions of n , where $\gamma^{(t)} = (\gamma_1^{(t)}, \gamma_2^{(t)}, \dots, \gamma_{p_t}^{(t)})$. We say that a (p_1, \dots, p_k) -colored factorization $(\pi_1, \dots, \pi_k, \phi_1, \dots, \phi_k)$ has *color-compositions* $(\gamma^{(1)}, \dots, \gamma^{(k)})$ if the permutation π_t has $\gamma_i^{(t)}$ elements colored i (i.e. $\gamma_i^{(t)} = |\phi_t^{-1}(i)|$) for all $t \in [k]$ and all $i \in [p_t]$. Let $c(\gamma^{(1)}, \dots, \gamma^{(k)})$ be the number of colored factorizations of color-compositions $(\gamma^{(1)}, \dots, \gamma^{(k)})$. In Section 4 we shall prove bijectively the following surprising symmetry property.

Theorem 1.3 (Symmetry property). *Let $\gamma^{(1)}, \delta^{(1)}, \dots, \gamma^{(k)}, \delta^{(k)}$ be compositions of n . If for every $t \in [k]$ the length of the compositions $\gamma^{(t)}$ and $\delta^{(t)}$ are equal, then $c(\gamma^{(1)}, \dots, \gamma^{(k)}) = c(\delta^{(1)}, \dots, \delta^{(k)})$.*

Given that there are $\binom{n-1}{\ell-1}$ compositions of n with ℓ parts, the symmetry property together with Theorem 1.2 gives the following refined formula.

Corollary 1.4. *For any compositions $\gamma^{(1)}, \dots, \gamma^{(k)}$ of n , the number of colored factorizations of color-compositions $(\gamma^{(1)}, \dots, \gamma^{(k)})$ is*

$$(1.3) \quad c(\gamma^{(1)}, \dots, \gamma^{(k)}) = \frac{n!^{k-1} M_{p_1-1, \dots, p_k-1}^{n-1}}{\prod_{t=1}^k \binom{n-1}{\ell(\gamma^{(t)})-1}},$$

Equation (1.3) is precisely the result established by Morales and Vassilieva for the cases $k = 2, 3$ in [10, 11]. These refined results can actually be obtained by a representation theory approach but we have not found them explicitly in the literature for $k > 3$.

As a final remark, observe that connection coefficients $\kappa(\lambda^{(1)}, \dots, \lambda^{(k)})$ are determined by the numbers $c(\lambda^{(1)}, \dots, \lambda^{(k)})$ of colored factorizations through a change of basis for symmetric functions:

$$\begin{aligned} \sum_{\lambda^{(1)}, \lambda^{(2)}, \dots, \lambda^{(k)}} \kappa(\lambda^{(1)}, \dots, \lambda^{(k)}) \prod_{t=1}^k p_{\lambda^{(t)}}(x_{t,1}, x_{t,2}, \dots) \\ = \sum_{\lambda^{(1)}, \lambda^{(2)}, \dots, \lambda^{(k)}} c(\lambda^{(1)}, \dots, \lambda^{(k)}) \prod_{t=1}^k m_{\lambda^{(t)}}(x_{t,1}, x_{t,2}, \dots), \end{aligned}$$

where the sums are over k -tuples of partitions of n and p_λ , m_λ denote respectively the *power sum* and *monomial* symmetric functions.

A probabilistic puzzle. We now describe a probabilistic puzzle associated to the set $\mathcal{M}_{p_1, \dots, p_k}^n$ appearing in Theorem 1.2.

Definition 1.5. For an integer t in $[k]$ and a subset $R \subsetneq [k]$ we define the integer $\alpha(t, R) \in [k]$ by setting

- $\alpha(t, R) = t - 1$ modulo k if $t \in R$,
- $\alpha(t) = t + r$ modulo k if $t \notin R$, $t + 1, \dots, t + r \in R$ and $t + r + 1 \notin R$.

Furthermore, for integers i_1, \dots, i_{k-1} in $[n]$ (repetitions allowed) and subsets $R_1, \dots, R_n \subsetneq [k]$, we define $\alpha((i_1, \dots, i_{k-1}), (R_1, \dots, R_n))$ to be the graph with vertex set $[k]$ and edge set $E = \{e_1, \dots, e_{k-1}\}$, where e_t is the edge $\{t, \alpha(t, R_{i_t})\}$.

In Section (6) we shall prove bijectively that Jackson's Theorem 1.2 is equivalent to the following probabilistic statement.

Theorem 1.6. Fix an integer $k \geq 2$ and positive integers n, p_1, \dots, p_k and consider the uniform distribution on pairs $B = ((i_1, \dots, i_{k-1}), (R_1, \dots, R_n))$ such that i_1, \dots, i_{k-1} are integers in $[n]$ (repetitions allowed) and $(R_1, \dots, R_n) \in \mathcal{M}_{p_1, \dots, p_k}^n$. Then the probability that the graph $\alpha(B)$ is a tree is equal to the probability that the subset R_1 has cardinality $k - 1$.

In Section 7, we shall give a direct proof of Theorem 1.6 for the cases $k = 2, 3, 4$ (thereby establishing Theorem 1.2 for these cases), and indicate how to handle a few more cases using a computer. In the forthcoming paper [3], we shall give a direct proof of Theorem 1.6 valid for all $k \geq 2$, and establish several other results of a similar flavor.

Constellations, or the drawings of factorizations. Our bijections are described in terms of k -constellations which are certain *maps* (embeddings of graphs in surfaces considered up to homeomorphism) with k types of vertices and with faces colored white and black (black faces are also called




FIGURE 2. Two hyperedge-labelled 3-constellations of size 5 (the shaded triangles represent the hyperedges). The 3-constellation on the left (which is embedded in the sphere) encodes the triple (π_1, π_2, π_3) , where $\pi_1 = (1, 2, 5)(3, 4)$, $\pi_2 = (1, 3)(2)(4)(5)$, $\pi_3 = (1, 4)(2)(3)(5)$, so that $\pi_1\pi_2\pi_3 = (1, 3, 2, 5)(4)$. The 3-cactus on the right (which is embedded in the torus) encodes the triple (π_1, π_2, π_3) , where $\pi_1 = (1, 3, 5)(2, 4)$, $\pi_2 = (1, 4)(2, 3)(5)$, $\pi_3 = (1)(2, 4)(3)(5)$, so that $\pi_1\pi_2\pi_3 = (1, 2, 3, 4, 5)$.

hyperedges and are colored in gray in our figures). We refer the reader to Section 2 for precise definitions, and to Figure 2 for some examples. To a k -constellation with its hyperedges labelled $1, \dots, n$, one associates the permutations π_1, \dots, π_k of $[n]$ where the cycles of the permutation π_t are in correspondence with the vertices of type t : the cycle associated to a vertex v is given by the counterclockwise order of the black faces incident to v ; see Figure 2. Actually, any tuple of permutations π_1, \dots, π_k of $[n]$ acting transitively on $[n]$ is associated to a unique k -constellation, thus k -constellations give a canonical way of “drawing” such tuples. Moreover, the number of cycles of the product $\pi_1 \circ \pi_2 \cdots \pi_k$ corresponds to the number of white faces of the associated constellation. In particular, the tuples (π_1, \dots, π_k) of permutations of $[n]$ such that the product $\pi_1 \circ \pi_2 \cdots \pi_k$ is an n -cycle (a permutation with 1 cycle) correspond bijectively to constellations with a single white face (because the transitivity condition is automatically satisfied in this case). Constellations with a single white face are called *cacti*.

Outline. In Section 2, we gather our definitions about maps, constellations and cacti. Theorems 1.2 and 1.3 can then be stated in terms of vertex-colored cacti.

In Section 3, we establish a bijection Φ between colored cacti and *tree-rooted constellations* which are constellations with a marked spanning tree. Through this bijection the number of colors of the vertices of type t in the cacti correspond to the number of vertices of type t in the corresponding tree-rooted constellation. This bijection is similar to the one used in [1] to prove the Harer-Zagier formula [5].

In Section 4, we establish a symmetry property for tree-rooted constellations which implies Theorem 1.3.

In Section 5, we decompose the *dual* of tree rooted maps and obtain certain decorated one-face maps that we call *nebulas*. Proving Jackson's counting formula then reduces to counting nebulas.

In Section 6 we give a bijection Ψ between nebulas and *valid biddings*, where the definition of bidding is closely related to the set $\mathcal{M}_{p_1, \dots, p_k}^n$ whose cardinality appears in Jackson's formula. We then prove that Jackson's formula is equivalent to the probabilistic statement given in Theorem 1.6.

Lastly, in Section 7 we prove the cases $k = 2, 3, 4$ of Theorem 1.6, thereby establishing Theorem 1.2 for these cases.

2. DEFINITIONS

For a positive integer n , we denote by $[n]$ the set $\{1, 2, \dots, n\}$. A *composition* of n is a sequence of positive integers $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_\ell)$ such that $\alpha_1 + \alpha_2 + \dots + \alpha_\ell = n$. The integers $\alpha_1, \dots, \alpha_\ell$ are the *parts* of α and the integer ℓ is the *length* of α . A *partition* is a composition $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_\ell)$ such that $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_\ell$.

Graphs and maps. Our *graphs* are undirected and can have loops and multiple edges. A *digraph*, or *directed graph*, is a graph where every edge is oriented; oriented edges are called *arcs*. An *Eulerian tour* of a directed graph is a directed path starting and ending at the same vertex and taking every arc exactly once. An edge e of a graph defines two *half-edges* each of them incident to an endpoint of e . A *rotation system* for a graph G is an assignment for each vertex v of G of a cyclic ordering for the half-edges incident to v .

We now review the connection between rotation systems and embeddings of graphs in surfaces. We call *surface* a compact, connected, orientable, 2-dimensional manifold without boundary (such a surface is characterized by its *genus* $g \geq 0$). A *map* is a cellular embedding of a connected graph in an oriented surface considered up to orientation preserving homeomorphism¹. By *cellular* we mean that the *faces* (connected components of the complement of the graph) are simply connected. For a map, the angular section between two consecutive half-edges around a vertex is called a *corner*. The *degree* of a vertex or a face is the number of incident corners. A map M naturally defines a rotation system $\rho(M)$ of the underlying graph G by taking the cyclic order of the half-edges incident to a vertex v to be the clockwise order of these half-edges around v . The following classical result (see e.g. [9]) states the relation between maps and graphs with rotation systems.

Lemma 2.1. *For any connected graph G , the function ρ is a bijection between the set of maps having underlying graph G and the set of rotation systems of G .*

¹Maps can be considered on non-orientable surfaces but we will not consider such surfaces here.

Constellations and cacti. A k -constellation, or *constellation* for short, is a map with two types of faces *black* and *white*, and k types of vertices $1, 2, \dots, k$, such that:

- (i) each edge separates a black face and a white face,
- (ii) each black face has degree k and is incident to vertices of type $1, 2, \dots, k$ in this order clockwise around the face.

Two constellations are shown in Figure 2. The black faces are also called *hyperedges*. The *size* of a constellation is the number of hyperedges. A constellation of size n is *labelled* if its hyperedges receive distinct labels in $[n]$.

We now recall the link between constellations and products of permutations. We call k -hypergraph a pair $G = (V, E)$ where V is a set of vertices, each of them having a *type* in $[k]$, and E is a set of *hyperedges* which are subsets of V containing exactly one vertex of each type. A *rotation-system* for the hypergraph G is an assignment for each vertex v of a cyclic order of the hyperedges *incident to* v (i.e., containing v). Clearly each k -constellation defines a connected k -hypergraph together with a rotation system (the clockwise order of the hyperedges around each vertex). In fact Lemma 2.1 readily implies the following result.

Lemma 2.2. *For any connected k -hypergraph G , there is a bijection between k -constellations of underlying k -hypergraph G and the rotation systems of G .*

Now given a hyperedge-labelled k -constellation C of size n , we define some permutations π_1, \dots, π_k as follows: for each $t \in [k]$ we define the cycles of the permutation π_t to be the counterclockwise order of the hyperedges around the vertices of type t . Examples are given in Figure 2. We then say that the hyperedge-labelled k -constellation C *represents* the tuple $\varrho(C) = (\pi_1, \dots, \pi_k)$. From Lemma 2.2 it is easy to establish the following classical result (see e.g. [7]).

Lemma 2.3. *The representation mapping ϱ is a bijection between hyperedge-labelled k -constellations of size n and tuples of permutations (π_1, \dots, π_k) of $[n]$ acting transitively on $[n]$. Moreover the number of white faces of the constellation is equal to the number of cycles of the product $\pi_1\pi_2 \cdots \pi_k$.*

An edge of a constellation has *type* $t \in [k]$ if its endpoints have types t and $t + 1$ (the types of the vertices and edges are considered modulo k). A k -constellation has *type* (p_1, \dots, p_k) if it has p_t vertices of type t for all $t \in [k]$. The *hyperdegree* of a vertex is the number of incident hyperedges. A constellation of type (p_1, \dots, p_k) is *vertex-labelled* if for each $t \in [k]$ the p_t vertices of type t have distinct labels in $[p_t]$. We say that such a constellation has *vertex-compositions* $(\gamma^{(1)}, \dots, \gamma^{(k)})$ if for all $t \in [k]$, $\gamma^{(t)}$ is a composition of size n and length p_t whose i th part is the hyperdegree of the vertex of type t labelled i .

A k -constellation is *rooted* if one of its hyperedges is distinguished as the *root hyperedge*. The vertex of type k incident to the root hyperedge is called

root vertex. There are $n!$ distinct ways of labelling a rooted constellation of size n (because a rooted constellation has no symmetry preserving the root hyperedge). Hence, there is a 1-to- $(n-1)!$ correspondence between rooted constellations of size n and hyperedge-labelled constellations of size n .

A k -*cactus* is a k -constellation with a single white face. By Lemma 2.3 the hyperedge-labelled k -cacti correspond bijectively to the factorizations of one of the $(n-1)!$ long cycles into k factors (transitivity is redundant in this case), while rooted cacti correspond bijectively to the factorizations of the permutation $(1, 2, \dots, n)$. Since Jackson's counting formula is about *colored factorizations* of $(1, 2, \dots, n)$ (see Definition 1.1), we now consider *vertex-colored cacti*. Given some positive integers q_1, \dots, q_k , a (q_1, \dots, q_k) -*colored cacti* is a k -cacti together with an assignment of *colors* to vertices, such that for every $t \in [k]$ the vertices of type t are colored using every color in $[q_t]$. A $(2, 1, 3)$ -colored cacti is represented in Figure 3. The *color-compositions* of a (q_1, \dots, q_k) -colored cacti of size n is the tuple $(\gamma^{(1)}, \dots, \gamma^{(k)})$, where for all $t \in [k]$, $\gamma^{(t)}$ is a composition of size n and length q_t whose i th part is the number of hyperedges incident to vertices of type t colored i . It is clear from the representation mapping ϱ , that (q_1, \dots, q_k) -colored cacti of color-compositions $(\gamma^{(1)}, \dots, \gamma^{(k)})$ are in bijection with the (q_1, \dots, q_k) -colored factorizations of $(1, 2, \dots, n)$ with color-compositions $(\gamma^{(1)}, \dots, \gamma^{(k)})$.




FIGURE 3. A $(2, 1, 3)$ -colored cacti (embedded in the sphere) with color-compositions $(\gamma^{(1)}, \gamma^{(2)}, \gamma^{(3)})$, where $\gamma^{(1)} = (1, 4)$, $\gamma^{(2)} = (5)$ and $\gamma^{(3)} = (2, 1, 2)$.

From now on, all our results and proofs are stated in terms of constellations and cacti.

3. FROM CACTI TO TREE-ROOTED CONSTELLATIONS

In this section we establish a bijection between vertex-colored cacti and certain constellations with a distinguished spanning tree. Let C be a k -constellation and let v_0 be a vertex. We call v_0 -*arborescence* of C a spanning tree A such that every vertex $v \neq v_0$ of type t is incident to exactly one edge of type t in A (equivalently, the spanning tree A is oriented from the leaves toward v_0 by orienting every edge of A of type $t \in [k]$ from its endpoint of type t toward its endpoint of type $t+1$). A *tree-rooted constellation* is a pair

(C, A) made of a rooted constellation C together with a v_0 -arborescence A , where v_0 is the root vertex of C . An example of tree-rooted constellation is given in Figure 4 (bottom right).

Theorem 3.1. *Let p_1, \dots, p_k be positive integers. There is a bijection Φ between the set $\mathcal{C}_{p_1, \dots, p_k}^n$ of (p_1, \dots, p_k) -colored rooted k -cacti of size n (these encode the (p_1, \dots, p_k) -colored factorizations of $(1, 2, \dots, n)$), and the set $\mathcal{T}_{p_1, \dots, p_k}^n$ of vertex-labelled tree-rooted k -constellations of size n and type (p_1, \dots, p_k) .*

Moreover, the bijection has the following degree preserving property: for any vertex-colored cactus C , the number of edges joining vertices of type t and color i to vertices of type $t+1$ and color j in C is equal to the number of edges joining the vertex of type t labelled i to the vertex of type $t+1$ labelled j in the tree-rooted constellation $\Phi(C)$.

Remark. The degree preserving property of Theorem 3.1 implies that for any tuple of compositions $(\gamma^{(1)}, \dots, \gamma^{(k)})$, the mapping Φ establishes a bijection between cacti of color-compositions $(\gamma^{(1)}, \dots, \gamma^{(k)})$ and tree-rooted constellations of vertex-compositions $(\gamma^{(1)}, \dots, \gamma^{(k)})$.

Remark. In the case $k = 2$ the tree-rooted k -constellations can be identified with rooted bipartite maps with a distinguished spanning tree (simply by considering the hyperedges as edges). These objects are easy to count (see [1]), so that the case $k = 2$ of Theorem 1.2 follows easily from Theorem 3.1 in this case.

The remaining of this section is devoted to the proof of Theorem 3.1. Our strategy parallels the one developed in [1] (building on some ideas of Lass [8]) in order to prove extensions of the Harer-Zagier formula. This proof is illustrated in Figure 4. We shall recombine the information given by a vertex-colored cactus into the information given by a tree-rooted constellation through the *BEST Theorem* (see Lemma 3.3 below).

We call k -digraph a directed graph with k types of vertices $1, \dots, k$, such that every vertex has as many ingoing and outgoing arcs, and every arc goes from a vertex of type t to a vertex of type $t+1$ for some $t \in [k]$ (as usual the types of vertices are considered modulo k). An arc going from a vertex of type t to a vertex of type $t+1$ is said to have *type* t . Note that a k -digraph has as many arcs of each type, and we say that it has *size* n if it has n arcs of each type. An *arc-labelling* of a k -digraph of size n is an assignment of distinct labels in $[n]$ to the n arcs of type t , in such a way that for any $(t, i) \in [k] \times [n]$ the end of the arc of type t and label i is the origin of the arc of type $t+1$ and label i . Observe that arc-labelled k -digraphs easily identify with hyperedge-labelled k -hypergraphs. A k -digraph has *type* (p_1, \dots, p_k) if for each $t \in [k]$ there are p_t vertices of type t . It is *vertex-labelled* by assigning distinct labels in $[p_t]$ to its p_t vertices of type t for all $t \in [k]$.




FIGURE 4. From a vertex-colored cactus to a tree-rooted constellation via the BEST Theorem.

Lemma 3.2. *There is a bijection Ξ between the set of hyperedge-labelled rooted (p_1, \dots, p_k) -colored cacti of size n , and the set of pairs (G, η) where G is a arc-labelled vertex-labelled k -digraph of type (p_1, \dots, p_k) and η is an Eulerian tour of G starting and ending at a vertex of type k .*

Lemma 3.2 is illustrated in the top part of Figure 4.

Proof. We call *black k -gon* a polygon with k vertices of type $1, 2, \dots, k$ in clockwise order, and *white kn -gon* a polygon with kn vertices, such that the type of vertices increases by one (modulo k) along each edge in counterclockwise order (modulo k). A white kn -gon is *rooted* if a corner incident to a vertex of type k is distinguished as the *root-corner*; it is (p_1, \dots, p_k) -colored if for all $t \in [k]$ the vertices of type t are colored using every color in $[p_t]$.

Observe that the n hyperedges of a k -cactus of size n are black k -gons, while its white face is a white kn -gon (since faces of cactus are simply connected). Moreover the k -cactus is completely determined (up to homeomorphism) by specifying the gluing of the black k -gons with the white kn -gon (that is specifying the pair of edges to be identified). Thus, a rooted hyperedge-labelled (p_1, \dots, p_k) -colored cactus is obtained by taking a rooted (p_1, \dots, p_k) -colored white kn -gon, and gluing its edges to the edges of n labelled black k -gon so as to respect the color and type of the vertices (certain vertices of the white kn -gon are identified by the gluing). Now, a rooted (p_1, \dots, p_k) -colored white kn -gon is bijectively encoded by a pair (\tilde{G}, η) , where \tilde{G} is a vertex-labelled k -digraph of type (p_1, \dots, p_k) and η is an Eulerian tour of \tilde{G} (the Eulerian tour gives the order of the colors around the white

kn -gon in counterclockwise direction starting from the root-corner). Moreover, the gluings of the n labelled black k -gons (respecting the type and coloring) are in bijection with the arc-labellings of \tilde{G} . This establishes the claimed bijection. \square

We now recall the BEST Theorem for Eulerian tours². Let G be a directed graph and let v_0 be a vertex. We call v_0 -Eulerian tour an Eulerian-tour starting and ending at vertex v_0 . Observe that a v_0 -Eulerian tour is completely characterized by its *local-order*, that is, the assignment for each vertex v of the order in which the outgoing edges incident to v are used. Note however that not every local order corresponds to an Eulerian tour. We call v_0 -arborescence a spanning tree A of G oriented from the leaves toward v_0 (i.e., every vertex $v \neq v_0$ has exactly one outgoing arc in A).

Lemma 3.3 (BEST Theorem). *Let G be an arc-labelled directed graph where every vertex has as many ingoing arcs as outgoing ones, and let v_0 be a vertex of G . A local order corresponds to a v_0 -Eulerian tour if and only if the set of last outgoing arcs out of the vertices $v \neq v_0$ form a v_0 -arborescence. Consequently, there is a bijection between the set of v_0 -Eulerian tours of G and the set of pairs (A, τ) , where A is a v_0 -arborescence, and τ is an assignment for each vertex v of a total order of the incident outgoing arcs not in A .*

We now complete the proof of Theorem 3.1. By combining Lemma 3.2 and the BEST Theorem, one gets a bijection between rooted hyperedge-labelled (p_1, \dots, p_k) -colored cacti and triples (G, A, θ) where G is an arc-labelled vertex-labelled k -digraph of type (p_1, \dots, p_k) , A is a v_0 -arborescence of G for a vertex v_0 of type k , and τ is an assignment for each vertex v of a total order of the arcs not in A going out of v . Observe that τ encodes the same information as a pair (a_0, τ') , where a_0 is an arc going out of v_0 and τ' is an assignment for each vertex v of a cyclic order of the arcs going out of v . Now the arc-labelled vertex-labelled k -digraph G encodes the same information as a hyperedge-labelled vertex-labelled k -hypergraph G' , and τ' can be seen as a rotation system for G' . Thus, by Lemma 2.2 the pair (G, τ) encodes the same information as a rooted hyperedge-labelled vertex-labelled k -constellation C of type (p_1, \dots, p_k) (note that the hypergraph G' is clearly connected since it has an arborescence A). Lastly, the v_0 -arborescence A of G clearly encodes a v_0 -arborescence of the constellation C , where v_0 is the root vertex of C .

We thus have obtained a bijection between rooted hyperedge-labelled (p_1, \dots, p_k) -colored cacti and the hyperedge-labelled vertex-labelled tree-rooted constellations. The labelling of the hyperedges can actually be disregarded since there are $n!$ distinct ways of labelling the hyperedges of a rooted constellation of size n . This gives the bijection announced in Theorem 3.1.

²This Theorem is due to de Bruijn, van Aardenne-Ehrenfest, Smith and Tutte. See [14, Theorem 5.6.2] for a proof.

Moreover it is easy to check that it has the claimed *degree preserving property*. \square

4. SYMMETRIES FOR TREE-ROOTED CONSTELLATIONS

In this section we prove that for vertex-labelled tree-rooted constellations of a given type (p_1, \dots, p_k) , every vertex-composition is equally likely. This together with Theorem 3.1 proves the symmetry property stated in Theorem 1.3.

We denote by $\mathcal{T}_{\gamma^{(1)}, \dots, \gamma^{(k)}}$ the set of vertex-labelled tree-rooted constellations of vertex-compositions $(\gamma^{(1)}, \dots, \gamma^{(k)})$.

Theorem 4.1. *If $\gamma^{(1)}, \dots, \gamma^{(k)}, \delta^{(1)}, \dots, \delta^{(k)}$ are compositions of n such that $\ell(\gamma^{(t)}) = \ell(\delta^{(t)})$ for all $t \in [k]$, then the sets $\mathcal{T}_{\gamma^{(1)}, \dots, \gamma^{(k)}}$ and $\mathcal{T}_{\delta^{(1)}, \dots, \delta^{(k)}}$ are in bijection.*

Remark. Theorem 4.1 gives the hope of counting tree-rooted constellations of given type, by looking at the simplest possible vertex-compositions. For instance, one can try to enumerate the set $\mathcal{T}_{\gamma^{(1)}, \dots, \gamma^{(k)}}$ where $\gamma^{(t)} = (n - p_t + 1, 1, 1, \dots, 1)$ for all $t \in [k]$ (similar ideas lead to a very easy way of counting k -cacti embedded in the sphere [2]). However, our efforts in this direction only led to a restatement of Jackson counting formula as a probabilistic puzzle similar to Theorem 1.6 which we could not easily solve for $k \geq 3$.

Proof. Let $t \in [k]$ and $i, j \in [p_t]$. In order to prove Theorem 4.1 it suffices to exhibit a bijection $\varphi_{t,i,j}$ between $\mathcal{T}_{\gamma^{(1)}, \dots, \gamma^{(k)}}$ and $\mathcal{T}_{\delta^{(1)}, \dots, \delta^{(k)}}$ when $\gamma^{(s)} = \delta^{(s)}$ for all $s \neq t$, $\gamma_x^{(t)} = \delta_x^{(t)}$ for all $x \neq i, j$, $\gamma_i^{(t)} - 1 = \delta_i^{(t)}$ and $\gamma_j^{(t)} + 1 = \delta_j^{(t)}$. In other words, we want to construct a bijection $\varphi_{t,i,j}$ which decreases by one the hyperdegree of the vertex of type t labelled i and increases by one the hyperdegree of the vertex of type t labelled j . Recall from Lemma 2.2 that a k -constellation is defined by a (connected) k -hypergraph together with a rotation system (clockwise order of hyperedges around the vertices); therefore it is well defined to *un glue* a hyperedge from a vertex of type t and *re glue* it in a specified corner of another vertex of type t . We will use these operations to define the mapping $\varphi_{t,i,j}$ below; see Figure 5.

Let $\mathcal{T}_{t,i}$ be the set of vertex-labelled tree-rooted constellations of type (p_1, \dots, p_k) such that the vertex of type t labelled i has hyperdegree at least two. Let T be a tree-rooted constellation in $\mathcal{T}_{t,i}$, let u_i and u_j be the vertices of type t labelled i and j respectively, let r be the root vertex, and let A be the marked r -arborescence. If $u_i \neq r$ we denote by h_i be hyperedge incident to the edge joining u_i to its parent in A , while if $u_i = r$ we denote by h_i the the root hyperedge. We define h_j similarly. Let h'_i be the hyperedge preceding h_i in clockwise order around u_i and let e_i be the edge of type $t-1$ incident to h'_i . Observe that $h_i \neq h'_i$ since the hyperdegree of u_i is at least two.

In order to define the mapping $\varphi_{t,i,j}$ we need to consider two cases which are illustrated in Figure 5. We first define a partition $\mathcal{T}_{t,i} = \mathcal{T}'_{t,i,j} \cup \mathcal{T}''_{t,i,j}$ by




FIGURE 5. The bijection $\varphi_{t,i,j}$ applied to a tree-rooted constellation in $\mathcal{T}'_{t,i,j}$ (left), or in $\mathcal{T}''_{t,i,j}$ (right). The tree-rooted k -constellations are represented as k -hypergraphs together with a rotation system (so the overlappings of the hyperedges in this figure are irrelevant).

declaring that T is in $\mathcal{T}'_{t,i,j}$ if the edge e_i is not on the path from u_j to the root vertex r in the arborescence A , and that T is in $\mathcal{T}''_{t,i,j}$ otherwise. Suppose first that T is in $\mathcal{T}'_{t,i,j}$. In this case we define $\varphi_{t,i,j}(T)$ as the constellation (with marked edges) obtained from the tree-rooted constellation T (with marked edges corresponding to the arborescence A) by ungluing the hyperedge h'_i from u_i and gluing it to u_j in the corner preceding the hyperedge h_j in clockwise order around u_j ; see Figure 5(a). Observe that $\varphi_{t,i,j}(T)$ is a tree-rooted constellation (in particular the marked edges form an r -arborescence A' of $\varphi_{t,i,j}(T)$). Moreover $\varphi_{t,i,j}(T)$ is in $\mathcal{T}_{t,j}$ and more precisely in $\mathcal{T}'_{t,j,i}$. It is also easy to see that $\varphi_{t,j,i}(\varphi_{t,i,j}(T)) = T$. Suppose now that T is in $\mathcal{T}''_{t,i,j}$. In this case we define $\varphi_{t,i,j}(T)$ as the constellation (with marked edges) obtained from T (with marked edges corresponding to the arborescence A) as follows: we unglue all the hyperedges incident to u_i except h_i and h'_i , we unglue all the hyperedges incident to u_j except h_j , we reglue the hyperedges unglued from u_j to u_i in the corner preceding h'_i in clockwise order around u_i (without changing their clockwise order), we reglue the hyperedges unglued from u_i to u_j (in the unique possible corner), and lastly we exchange the labels i and j of the vertices u_i and u_j ; see Figure 5(b). It is easy to see that $\varphi_{t,i,j}(T)$ is a tree-rooted constellation (in particular the marked edges form an r -arborescence of $\varphi_{t,i,j}(T)$). Moreover $\varphi_{t,i,j}(T)$ is in $\mathcal{T}_{t,j}$ and more precisely in $\mathcal{T}''_{t,j,i}$. It is also easy to see that $\varphi_{t,j,i}(\varphi_{t,i,j}(T)) = T$.

We have shown that $\varphi_{t,i,j}$ is a mapping from $\mathcal{T}_{t,i}$ to $\mathcal{T}_{t,j}$. Moreover $\varphi_{t,j,i} \circ \varphi_{t,i,j} = Id$ for all i, j , thus $\varphi_{t,i,j} = \varphi_{t,j,i}^{-1}$ is a bijection. Lastly, the bijection $\varphi_{t,i,j}$ decreases by one the hyperdegree of the vertex of type t labelled i and increases by one the degree of the vertex of type t labelled j . Thus $\varphi_{t,i,j}$ has all the claimed properties. \square

5. FROM TREE-ROOTED CONSTELLATIONS TO NEBULAS

In Section 3 we obtained a bijection between vertex-colored cacti and tree-rooted constellations. In this section we take another look at tree-rooted constellations by characterizing their *duals*. We eventually obtain a bijection between a class of tree-rooted constellations and certain decorated maps with a single face called *nebulas*.

We call *tree-pointed k -constellation* a pair (C, A) , where C is a rooted k -constellation and A is a v_0 -arborescence for some vertex v_0 which can be distinct from the the root-vertex of C . If the constellation C has type (p_1, \dots, p_k) and v_0 has type t , then the tree-pointed k -constellation (C, A) is said to have *reduced type* (p'_1, \dots, p'_k) , where $p'_t = p_t - 1$ and $p'_s = p_s$ for $s \neq t$. Observe that the arborescence A has p'_t edges of type t for all $t \in [k]$ (because for every vertex $v \neq v_0$ of type t the edge of A joining v to its parent has type t).

Lemma 5.1. *There is a 1-to- $\prod_{i=1}^k p_i!$ correspondence between tree-pointed k -constellations of reduced type (p_1, \dots, p_k) , and the union $\mathcal{T}_{p_1+1, p_2, \dots, p_k}^n \cup \mathcal{T}_{p_1, p_2+1, \dots, p_k}^n \cup \dots \cup \mathcal{T}_{p_1, p_2, \dots, p_k+1}^n$ of vertex-labelled tree-rooted constellations.*

Proof. We first claim that for any constellation C and any vertices u, v , there are as many u -arborescences as v -arborescences. Indeed, if one orients the edges of type t of C toward their endpoint of type $t + 1$, one gets a *Eulerian digraph* (oriented graph with as many ingoing and outgoing edges at each vertex). Moreover it is an easy corollary of the BEST Theorem that Eulerian digraphs have the same number arborescences directed toward each vertex (see [14, Cor 5.6.3]). Thus, the tree-pointed constellations of reduced type (p_1, \dots, p_k) such that the root vertex of the arborescence has type t are equinumerous to the tree-rooted constellations of type $(p_1, \dots, p_t + 1, \dots, p_k)$ with a distinguished vertex of type t . Since there are $(p_t + 1)$ ways of distinguishing a vertex of type t in such a tree-rooted constellation versus $(p_t + 1) \prod_{i=1}^k p_i!$ ways of labelling its vertices, one gets the claimed correspondence. \square

We now consider the dual of constellations. Recall that the *dual* of a map M is the map M^* obtained by placing a vertex of M^* in each face of M and drawing an edge of M^* across each edge of M . Duality is a genus preserving involution on maps such that the vertices, edges and faces of M correspond respectively to the faces, edges and vertices of M^* . We now describe the dual of k -constellations. Observe that k -constellations can be characterized as the maps with black and white faces, and k types of edges $1, 2, \dots, k$ such that

- (i) each edge separates a black and a white face,
- (ii) each black face has degree k ,
- (iii) in clockwise order around a hyperedge (resp. white face) the type of the edges increases (resp. decreases) by one from one edge to the next.

Indeed, with the preceding conditions, for each vertex v there exists $t \in [k]$ such that the edges incident to v are alternatingly of type $t - 1$ and t for some t in $[k]$ (and we can thus say that v has type t in this case). With the previous characterization, it is clear that duality gives a bijection between k -constellations and the *dual-constellations*, which are defined as the maps with black and white vertices, and k types of edges $1, 2, \dots, k$ such that

- (i) each edge joins a black and a white vertex
- (ii) each black vertex has degree k ,
- (iii) in clockwise order around a black vertex (resp. white vertex) the type of the edges increases (resp. decreases) by one from one edge to the next.

We will now describe a class of maps closely related to the dual of tree-rooted constellations. A *bud* is a dangling half-edge, that is, a half-edge which is not part of a complete edge. We define a k -*nebula*, or *nebula* for short, as a map having a *single face* with black and white vertices, k types of edges $1, \dots, k$, and k types of buds $1, \dots, k$ such that

- (i) each edge joins a black and a white vertex,
- (ii) each black vertex has degree k ,
- (iii) in clockwise order around a black vertex (resp. white vertex) the type of the edges or buds increases (resp. decreases) by one from one half-edge to the next,
- (iv) for all $t \in [k]$ the number of buds of type t incident to black vertices is equal to the number of buds of type t incident to white vertices.

A 3-nebula is shown in Figure 6 (right). A nebula is *rooted* if one of the black vertices is distinguished as the *root vertex*. We call *black buds* and *white buds* respectively the buds incident to black and to white vertices. A k -nebula is said to have *size* n and *type* (p_1, \dots, p_k) if it has n black vertices and p_t black buds of type t for all $t \in [k]$.

Consider a k -constellation C and a spanning tree A . We call *dual-opening* of (C, A) the map with buds N obtained from the dual-constellation C^* by cutting in two halves the edges of C^* which are crossing the edges of the tree A (the edges of C^* of type t crossing A gives two buds of type t). We call *root vertex* of N the dual of the root hyperedge of C . The dual-opening of a tree-pointed constellation (C, A) is represented in Figure 6.

We now state the main result of this section.

Theorem 5.2. *The dual-opening gives a bijection between tree-pointed constellations of reduced type (p_1, \dots, p_k) and rooted nebulas of type (p_1, \dots, p_k) .*

Theorems 3.1, 5.2 and Lemma 5.1 immediately imply the following result.

Corollary 5.3. *The nebulas of size n and type (p_1, \dots, p_k) are in 1-to- $\prod_{i=1}^k p_i!$ correspondence with the disjoint union $\mathcal{C}_{p_1+1, p_2, \dots, p_k}^n \cup \mathcal{C}_{p_1, p_2+1, \dots, p_k}^n \cup \dots \cup \mathcal{C}_{p_1, p_2, \dots, p_k+1}^n$ of colored cacti.*




FIGURE 6. Dual-opening of a tree-pointed constellation. Here the tree-constellation is embedded in the sphere (hence so is its dual opening). The types of edges are indicated by their thicknesses. The root-hyperedge is indicated in darker shade, and the corresponding root-vertex of the nebula is represented in gray.

Remark. By definition, tree-rooted constellations are tree-pointed constellations such that the marked arborescence is directed toward the root-vertex of the map. One can check that tree-rooted constellations actually correspond to nebulas such that the sequence of black buds and white buds around N form a *parenthesis system*, that is, the number of black buds never exceeds the number of white buds when turning in clockwise direction around the face of N starting at the corner preceding the edge of type 1 around the root vertex. Our reason for considering tree-pointed constellations is to get rid of the parenthesis system condition (which is hard to control).

The rest of this section is devoted to the proof of Theorem 5.2.

Lemma 5.4. *Let C be a k -constellation and let A be a spanning tree (not necessarily an arborescence) having p_t edges of type $t \in [k]$. Then the dual-opening of (C, A) is a nebula of type (p_1, \dots, p_k) .*

Proof. Let N be the dual opening of (C, A) . Since the spanning tree A connects all the vertices of the constellation C , the faces of C^* are all merged into a single face of N by cutting the edges of C^* crossed by A . Moreover, this face of N is simply connected because A is simply connected (i.e., has no cycle). Hence N is indeed a *map* with a *single face*. The other properties of nebulas are easily seen to hold. \square

Lemma 5.4 proves that the dual opening of a tree-pointed constellation is a nebula. We now define the closure of nebulas. Let N be a nebula. We call *turning clockwise around N* the process of walking around the face of N by following its edges, with the edges on the right side of the walker (see Figure 8(a)). We say that a white bud w *matches* a black bud b if there is no bud between w and b when turning clockwise around N starting from w . We call *closure* of the nebula N , the result of *recursively* forming edges by gluing together pairs of matching buds (thus, at a later step of this recursive

process, we say that a white bud w matches a black bud b , if there is no bud between w and b when turning clockwise around the face containing the buds). This process is illustrated in Figure 7. It is clear that the closure can be done without creating any edge-crossings, that it will exhaust all buds, and that the result is uniquely defined.

Lemma 5.5. *The pairs of buds glued together during the closure of a nebula have the same type. Consequently, the closure of a nebula gives a dual-constellation.*

Proof. It is easy to see that if a map satisfies Condition (iii) of nebulas, then matching pairs of buds (w, b) have the same type. Moreover, in this case, Condition (iii) is preserved by forming an edge out of the buds w, b . This shows the first claim by induction on the number of buds. The second claim is clear. \square




FIGURE 7. The closure of a nebula N (here N is embedded in the sphere) obtained by recursively gluing pairs of matching edges gives a dual-constellation D (with dashed bud-edges). Taking its dual gives a tree-rooted constellation (C, A) .

Let N be a nebula, let D be its closure (which is a dual-constellation) and let C be the corresponding constellation. Let A be the set of edges of the constellation C which are dual to the edges of D which have been created during the closure of N (by joining two buds). The pair (C, A) is called the *dual-closure* of the nebula N .

Lemma 5.6. *The dual-closure of a nebula is a tree-pointed constellation.*

Proof. Let N be a nebula, let D be its closure, and let (C, A) be its dual-closure. By Lemma 5.5 we know that C is a constellation, so we only need to show that A is an arborescence of C . We call bud-edges the edges of D created during the closure and we view them as oriented from the white bud to the black bud (so that we can distinguish the face on their *left* and the face on their *right*). Let f_0 be the face of the nebula N . During the closure

of N , each time a matching pair of buds are glued into a bud-edge e , we consider the face at the right of e as a “new face”, while we consider the face the left of e as the “original face” f_0 . Hence, at any time during the closure, the original face f_0 is at the left of every incident bud-edge, while any face $f \neq f_0$ is at the right of exactly one bud-edge which we call the *closing edge of f* . Observe that for any face $f \neq f_0$ of D , one can reach the face f_0 of D by starting inside f and repeatedly crossing the closing edge of the current face.

Let v_0 be the vertex of the constellation C dual to the face f_0 of D , and for any vertex $v \neq v_0$ let the *parent edge* of v be the dual of the closing edge of f , where f is the face of D dual to the vertex v . By definition, A is the set of parent edges and we will show that it is a v_0 arborescence. First of all, the above observation implies that starting from a vertex $v \neq v_0$ and repeatedly following the parent edges one eventually reaches the vertex v_0 . Thus A connects all the vertices of C and since there is one more vertex in C than edges in A , we can conclude that A is a spanning tree of C . It now suffices to check that the tree A is oriented toward v_0 , or equivalently that the parent edge of a vertex of type t has type t . This is true because of the orientation convention for closing edges and the fact that the type increases in clockwise order around the black vertices of the dual-constellations D . \square

We now conclude the proof of Theorem 5.2. Let Λ and Δ denote respectively the dual-opening and dual-closure mappings. By Lemma 5.6, the closure $\Delta(N)$ of any nebula N is a tree-pointed constellation. Moreover, it is clear that $\Lambda \circ \Delta(N) = N$.

We now consider a tree-pointed constellation (C, A) , where A denotes the marked v_0 -arborescence. By Lemma 5.4, $N := \Lambda(C, A)$ is a nebula. We want to prove that the dual-closure of N is the tree-pointed constellation (C, A) . For this, we consider the white bud w and black bud b obtained by cutting in two halves an edge e^* of the dual-constellation C^* crossing A , and want to prove that the buds w and b will be glued together during the closure of N . We consider the set of buds encountered between w and b when turning clockwise around N starting from the bud w . Note that when N and A are superimposed, turning around the face of N is the same as turning (counterclockwise) around the arborescence A . Let e be the edge of A crossing the edge e^* , and let A' be the subtree of $A - \{e\}$ not containing the vertex v_0 . The buds between w and b around N are the buds cut by the subtree A' (since the tree A is oriented toward v_0). We now reason by induction on the number of edges in A' to show that w and b will be glued together during the closure of N . First observe that if A' has no edge, then w and b are matching buds, hence they will indeed be glued together during the closure of N . Now if A' has some edges, we know by induction that all the buds between w and b will be glued in pairs during the closure, therefore w and b will eventually be matching, hence will be glued together during the

closure of N . Thus, the dual-closure of N is the tree-pointed constellation (C, A) . In other words, $\Delta \circ \Lambda(C, A) = (C, A)$.

We have proved that Λ and Δ are inverse mappings, so the dual-opening Λ is a bijection. This completes the proof of Theorem 5.2. \square

6. FROM NEBULAS TO BIDDINGS

In this section we draw a connection between nebulas and the set $\mathcal{M}_{p_1, \dots, p_k}^n$ whose cardinality appears in Theorem 1.2. More precisely we will encode nebulas by *biddings*, where a *bidding* of size n and type (p_1, \dots, p_k) is a pair $((\omega_1, \dots, \omega_k), (R_1, \dots, R_n))$, where $\omega_1, \dots, \omega_k$ are permutations of $[n]$, and (R_1, \dots, R_n) belongs to $\mathcal{M}_{p_1, \dots, p_k}^n$. This will lead to the probabilistic puzzle stated in Theorem 1.6.

A nebula of size n and type (p_1, \dots, p_k) is said to be *labelled* if its n black vertices are given distinct labels in $[n]$, and the p_t white buds of type t are given distinct labels in L_t , where L_t is the set of labels of the p_t black vertices incident to the black buds of type t . Clearly there are $n! \prod_t p_t!$ ways of labelling a nebula of size n and type (p_1, \dots, p_k) . We also define the label of any edge of the nebula to be the label of the incident black vertex, so that for every pair $(t, i) \in [k] \times [n]$ there is either an edge or a white bud of type t and label i . We denote by $c(t, i)$ the corner preceding either the edge or white bud of type t labelled i in clockwise order around the incident white vertex.

Let N be a labelled rooted k -nebula of size n . We will now consider the sequence of *white corners* (corners incident to white vertices) encountered when turning around the nebula N . Recall that *turning clockwise around N* means walking around the face of N by following the edges of N , with the edges on the right-side of the walker (the buds are just *crossed*). A clockwise tour of a nebula is indicated in Figure 8. By turning clockwise around N each of the white corners $\{c(t, i), t \in [k], i \in [n]\}$ are visited, and this defines a cyclic order on $[k] \times [n]$. This cyclic order then gives a total order on the set $[k] \times [n]$ by choosing (k, ℓ_0) to be the greatest element, where ℓ_0 is the label of the (black) root vertex of N . We call this total order the *appearance order* on $[k] \times [n]$ and we denote it by \prec_N ,

Definition 6.1. *For a rooted nebula N , we denote by $\Psi(N)$ the bidding $((\omega_1, \dots, \omega_k), (R_1, \dots, R_n))$ defined as follows:*

- (i) *for all $i \in [n]$, $R_i \subsetneq [k]$ is the set of types of the buds incident to the black vertex labelled i ,*
- (ii) *for all $t \in [k]$, ω_t is the permutation of n giving the appearance order of the edges or white buds of type t , that is, $(t, \omega_t(1)) \prec_N (t, \omega_t(2)) \prec_N \dots \prec_N (t, \omega_t(n))$.*

The mapping Ψ is indicated in Figure 8. Recall now Definition 1.5 of the mapping α . We say that a bidding $((\omega_1, \dots, \omega_k), (R_1, \dots, R_n))$ is *valid* if

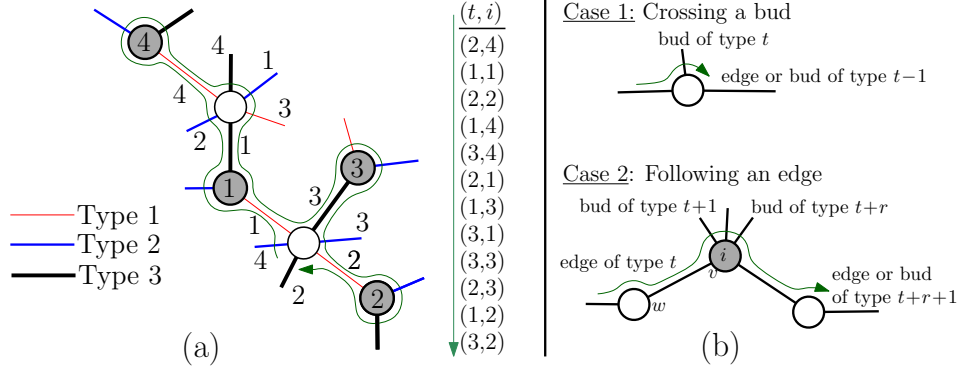


FIGURE 8. (a) The clockwise tour of a nebula N . Here $k = 3$, $n = 4$ and the root vertex has label $\ell_0 = 2$. The appearance order on $[k] \times [n]$ is indicated vertically (from top to bottom). The bidding $\Psi(N) = ((\omega_1, \omega_2, \omega_3), (R_1, R_2, R_3, R_4))$ is given by $R_1 = \{2\}$, $R_2 = \{2, 3\}$, $R_3 = \{1, 2\}$, $R_4 = \{2, 3\}$ and $\omega_1 = 1432$, $\omega_2 = 3214$, $\omega_3 = 4132$ (the permutations are indicated in one-line notation here). (b) From a white corner to the next during the clockwise tour of a nebula.

the graph $\alpha((\omega_1(n), \dots, \omega_{k-1}(n)), (R_1, \dots, R_n))$ is a tree. We now state the main result of this section.

Theorem 6.2. *The mapping Ψ is a bijection between labelled rooted nebulas of size n and type (p_1, \dots, p_k) and valid biddings of size n and type (p_1, \dots, p_k) .*

Corollary 6.3. *The valid biddings of size n and type (p_1, \dots, p_k) are in $n!$ -to-1 correspondence with the disjoint union $\mathcal{C}_{p_1+1, p_2, \dots, p_k}^n \cup \dots \cup \mathcal{C}_{p_1, p_2, \dots, p_k+1}^n$ of colored cacti. Consequently, Jackson's counting formula (Theorem 1.2) is equivalent to Theorem 1.6.*

Proof of Corollary 6.3. Recall that there are $n! \prod_t p_t!$ ways of labelling a rooted nebulas of size n and type (p_1, \dots, p_k) . Hence, by Corollary 5.3, labelled rooted nebulas of size n and type (p_1, \dots, p_k) are in $n!$ -to-1 correspondence with the union $\mathcal{C}_{p_1+1, p_2, \dots, p_k}^n \cup \dots \cup \mathcal{C}_{p_1, p_2, \dots, p_k+1}^n$ of colored cacti. Thus the first statement of Corollary 6.3 is a direct consequence of Theorem 6.2.

We now prove the second statement of Corollary 6.3. First observe that Theorem 1.2 implies that the set $\tilde{\mathcal{C}}_{p_1, p_2, \dots, p_k}^n := \mathcal{C}_{p_1+1, p_2, \dots, p_k}^n \cup \dots \cup \mathcal{C}_{p_1, p_2, \dots, p_k+1}^n$ has cardinality

$$|\tilde{\mathcal{C}}_{p_1, p_2, \dots, p_k}^n| = n!^{k-1} \left(M_{p_1, p_2-1, \dots, p_k-1}^{n-1} + \dots + M_{p_1-1, \dots, p_{k-1}-1, p_k}^{n-1} \right).$$

Conversely it is easy to see that this equation implies Theorem 1.2 (that is, $|\mathcal{C}_{p_1, p_2, \dots, p_k}^n| = n!^{k-1} M_{p_1-1, p_2-1, \dots, p_k-1}^{n-1}$) by induction on p_1 starting with the

base case

$$(6.1) \quad |\mathcal{C}_{1,p_2,\dots,p_k}^n| = n!^{k-1} \prod_{t=2}^k \binom{n-1}{p_t-1} = n!^{k-1} M_{0,p_2-1,\dots,p_k-1}^{n-1}.$$

This base case can be checked as follows. The first equality in (6.1) holds because there are $n! \binom{n-1}{p-1}$ permutations with cycles colored using every color in $[p]$, and for any permutations π_2, \dots, π_k there is a unique permutation π_1 such that $\pi_1 \circ \pi_2 \circ \dots \circ \pi_k = (1, 2, \dots, n)$ and a unique coloring of the cycles of this permutation π_1 with 1 color. The second equality in (6.1) holds because $M_{0,p_2-1,\dots,p_k-1}^{n-1}$ is the set of $(n-1)$ -tuples of subsets of $\{2, \dots, k\}$ for which every $t \in \{2, \dots, k\}$ appears in exactly $p_t - 1$ subsets (so that there are $\binom{n-1}{p_t-1}$ ways of choosing in which subsets t appears).

Observe now that $M_{p_1,p_2-1,\dots,p_k-1}^{n-1} + \dots + M_{p_1-1,\dots,p_k-1-1,p_k}^{n-1}$ can be interpreted as the cardinality of the set $\widetilde{\mathcal{M}}_{p_1,p_2,\dots,p_k}^n$ of tuples (R_1, \dots, R_n) in $\mathcal{M}_{p_1,p_2,\dots,p_k}^n$ such that R_1 has cardinality $k-1$ (indeed $M_{p_1,p_2-1,\dots,p_k-1}^{n-1}$ counts the tuples (R_1, \dots, R_n) in $\mathcal{M}_{p_1,p_2,\dots,p_k}^n$ such that $R_1 = \{2, 3, \dots, k\}$ etc.). Therefore the preceding discussion shows that Theorem 1.2 is equivalent to

$$(6.2) \quad n! |\widetilde{\mathcal{C}}_{p_1,p_2,\dots,p_k}^n| = n!^k |\widetilde{\mathcal{M}}_{p_1,p_2,\dots,p_k}^n|.$$

Now, by Theorem 6.2, the left-hand side of (6.2) is the number of valid biddings of size n and type (p_1, \dots, p_k) , while the right-hand side is the number of biddings of size n and type (p_1, \dots, p_k) such that R_1 has cardinality $k-1$. This shows that Theorem 1.2 is equivalent to Theorem 1.6. \square

The remaining of this section is devoted to the proof of Theorem 6.2. In order to analyze the mapping Ψ we introduce an intermediate class of objects called *prebidding* and define a mapping ϑ between labelled nebulas and *prebiddings*, and then a mapping σ between *prebiddings* and biddings such that $\Psi = \sigma \circ \vartheta$.

A *prebidding* is a pair $(\prec, (R_1, \dots, R_n))$ where \prec is a linear order on $[k] \times [n]$ and R_1, \dots, R_n are strict subsets of $[k]$. We say that a prebidding is *valid* if the greatest element is of the form (k, ℓ_0) for some $\ell_0 \in [n]$ and if, whenever (t, i) and (t', i') are consecutive pairs in the order \prec , or when $(t, i) = (k, \ell_0)$ and (t', i') is the least element, one has $t' = \alpha(t, R_i)$ (that is, $t' = t - 1$ if $t \in R_i$ and $t' = t + r$ if $t \notin R_i$, $t + 1, \dots, t + r \in R_i$ and $t + r + 1 \notin R_i$).

We now consider the mapping ϑ which associates to a labelled rooted k -nebula N the prebidding $\vartheta(N) = (\prec_N, (R_1, \dots, R_n))$, where \prec_N is the appearance order on $[k] \times [n]$ and $R_i \subsetneq [k]$ is the set of types of the buds incident to the black vertex labelled i .

Lemma 6.4. *The mapping ϑ is a bijection between labelled rooted nebulas and valid prebiddings.*

Proof. Let N be a nebula. We first prove that the prebidding $\vartheta(N) = (\prec_N, (R_1, \dots, R_n))$ is valid. It is clear from the definition of the appearance order \prec_N that the greatest element is of the form (k, ℓ_0) for some $\ell_0 \in [n]$. We now consider pairs (t, i) and (t', i') which are either consecutive in the appearance order or such that $(t, i) = (k, \ell_0)$ and (t', i') is the least element of the appearance order. By definition, if $t \in R_i$ the pair (t, i) corresponds to a white bud of N and if $t \notin R_i$ it corresponds to an edge of N . We first suppose that $t \in R_i$. Since $t \in R_i$, the white corner $c(t, i)$ precedes the white bud b of type t labelled i , so that after visiting the corner $c(t, i)$ the clockwise tour of N crosses the bud b and arrives at a white corner preceding either a bud or an edge of type $t - 1$ (see Figure 8(b)). Thus $t' = t - 1$. Suppose now that $t \notin R_i$. In this case the pair (t, i) corresponds to an edge e incident to the black vertex v labelled i and to a white vertex w . After passing through the white corner $c(t, i)$ preceding e around w , the clockwise tour follows e , crosses the buds of type $t + 1, \dots, t + r \in R_i$ around the black vertex v , then follows an edge of type $t + r + 1 \notin R_i$ and arrives at a white corner preceding either an edge or a bud of type $t + r$ (see Figure 8(b)). Thus $t' = t + r$. This completes the proof that the prebidding $\vartheta(N)$ is valid.

We now prove the injectivity of the mapping ϑ . Observe that from the prebidding $\vartheta(N)$ one can deduce the complete list of edges followed and buds crossed during the clockwise tour of the nebula N . This determines a polygon with black and white vertices, and with buds drawn on the inside region. This is illustrated in Figure 9. Moreover the edges come in pairs having the same types and labels so that one knows how to glue the edges of the polygon in pairs in order to recover the nebula N . Lastly the root vertex of N is identified as the black vertex of label ℓ_0 , where (k, ℓ_0) is the greatest element of \prec_N .

In order to prove the surjectivity of ϑ we consider a valid prebidding $P = (\prec, (R_1, \dots, R_n))$ and want to check that one can apply the above mentioned procedure to get a nebula (see Figure 9). From the prebidding P one can construct a polygon with black and white vertices, and with buds drawn on the inside region such that edges and buds have types which increase clockwise around white vertices and decrease clockwise around white vertices (the polygon is constructed in such a way that its clockwise tour is described by the prebidding P). Now it is easy to check that for each type t and label i either there is both a white bud and a black bud of type t labelled i (this happens if $t \in R_i$) or there are two edges of type t labelled i one going from a black vertex to a white vertex, and one going the opposite way when turning clockwise around the polygon (this happens if $t \notin R_i$). Thus one can glue together the pairs of edges of the same types and labels, and thereby obtain a map N with a single face such that its clockwise tour is described by the prebidding P . In the map N the black vertices have degree k since for every $i \in [n]$ the black vertex labelled i is incident exactly to the edges and black buds labelled i . Thus N is a nebula. Lastly it is clear that $\vartheta(N) = P$. This proves the surjectivity and hence the bijectivity of ϑ . \square




FIGURE 9. From valid prebiddings to nebulas (the appearance order is indicated from top to bottom).

We now define a mapping σ on valid prebiddings by setting

$$\sigma(\prec, (R_1, \dots, R_n)) = ((\omega_1, \dots, \omega_k), (R_1, \dots, R_n)),$$

where ω_t is the permutation of $[n]$ defined by $(t, \omega_t(1)) \prec (t, \omega_t(2)) \prec \dots \prec (t, \omega_t(n))$. Observe that $\Psi = \sigma \circ \vartheta$. The proof of the following lemma is based on the BEST Theorem.

Lemma 6.5. *The mapping σ is a bijection between valid prebiddings and valid biddings.*

Proof. Let (R_1, \dots, R_n) be a fixed tuple of strict subsets of $[k]$. We consider the directed graph G with vertex set $[k]$ and arc set $\{a_{t,i}, (t,i) \in [k] \times [n]\}$, where $a_{t,i}$ is an arc from vertex t to vertex $\alpha(t, R_i)$. First observe that every vertex t of the graph G has n outgoing arcs (the arcs $a_{t,i}$, $i \in [n]$) and n ingoing arcs (because for each $i \in [n]$ there a unique $t' \in [k]$ such that $\alpha(t', R_i) = t$). Now, to any Eulerian tour of G one can associate a linear order \prec on $[k] \times [n]$ defined by setting $(t,i) \prec (t',i')$ if the arc $a_{t,i}$ is taken before the arc $a_{t',i'}$ during the Eulerian tour. It is easily seen that this gives a bijection σ_1 between the linear orders \prec such that $(\prec, (R_1, \dots, R_n))$ is a valid prebidding and the Eulerian tours of G starting (and ending) at vertex k . Moreover, by the BEST Theorem, the Eulerian tours of G starting at vertex k are in bijection with the assignments for each vertex $t \in [k]$ of a linear order of the arcs going out of t (the order in which these arcs are used during the Eulerian tour) such that the set of greatest arcs going out of the vertices $1, 2, \dots, k-1$ form a spanning tree of G directed toward k . Note that the linear order on the outgoing edges of a vertex t can be represented by a permutation ω_t defined by setting $\omega_t(i) = j$ if $a_{t,j}$ is the arc going out of vertex t used at the i th exit of that vertex during the Eulerian tour. Hence the BEST theorem gives a bijection σ_2 between the Eulerian tours of G and the tuples of permutations $(\omega_1, \dots, \omega_k)$ such that $((\omega_1, \dots, \omega_k), (R_1, \dots, R_n))$ is a valid bidding. This completes the proof since $\sigma(\prec, (R_1, \dots, R_n)) = (\sigma_2 \circ \sigma_1(\prec), (R_1, \dots, R_n))$. \square

By definition $\Psi = \sigma \circ \vartheta$, so Theorem 6.2 is a direct consequence of Lemmas 6.4 and 6.5. \square

7. CASES $k = 2, 3, 4$ OF THEOREM 1.6.

We have shown (Corollary 6.3) that for each $k \geq 2$ Jackson's counting formula is equivalent to Theorem 1.6. In this section we prove a few cases of Theorem 1.6.

Case $k = 2$. Let $B = ((i), (R_1, \dots, R_n))$, where $i \in [n]$ and $(R_1, \dots, R_n) \in \mathcal{M}_{p_1, p_2}^n$. The graph $G = \alpha(B)$ has vertex set $\{1, 2\}$ and one arc $a_1 = (1, x)$ where $x = 2$ if $R_i = \{1\}$ or $R_i = \{2\}$, and $x = 1$ otherwise (if R_i is empty). Thus, G is a tree if and only if R_i has cardinality 1. Clearly the probability that R_i has cardinality 1 is the same as the probability that R_1 has cardinality 1, thus Theorem 1.6 is proved for $k = 2$.

Case $k = 3$. Let $B = ((i, j), (R_1, \dots, R_n))$, where $i, j \in [n]$ and $(R_1, \dots, R_n) \in \mathcal{M}_{p_1, p_2, p_3}^n$. The graph $G = \alpha(B)$ has vertex set $\{1, 2, 3\}$ and two arcs $a_1 = (1, x)$, and $a_2 = (2, y)$. We want to evaluate the probability that the graph G is a tree. There are three possible trees and they occur in the events described by Figure 10. Using inclusion-exclusion one gets that the probability that G is a tree is

$$(7.1) \quad \begin{aligned} \mathbb{P}(\text{tree}) = & P_{\{1\} \times \{2\}} + P_{\{1\} \times \{3\}} + P_{\{2\} \times \{3\}} \\ & - P_{\{1\} \times \{2, 3\}} - P_{\{2\} \times \{1, 3\}} - P_{\{3\} \times \{1, 2\}}, \\ & + P_{\{1, 2\} \times \{1, 3\}} + P_{\{1, 2\} \times \{2, 3\}} + P_{\{1, 3\} \times \{2, 3\}}, \end{aligned}$$

where $P_{A \times B}$ denotes the probability of the event $\{A \subseteq R_i \text{ and } B \subseteq R_j\}$ (all these probabilities depend on n, p_1, p_2, p_3). We then use the following easy lemma.

Lemma 7.1 (Exchange Lemma for $k = 3$). *For $\{a, b, c\} = \{1, 2, 3\}$ one has*

$$P_{\{a, b\} \times \emptyset} = P_{\{a\} \times \{b\}} - P_{\{a, c\} \times \{b\}} + P_{\{a, b\} \times \{a, c\}}.$$

Proof. We consider the events $E_1 = \{a \in R_i, c \notin R_i \text{ and } b \in R_j\}$ and $E_2 = \{a, b \in R_i, c \notin R_i \text{ and } R_j \neq \{a, c\}\}$. The probabilities of these event are $\mathbb{P}(E_1) = P_{\{a\} \times \{b\}} - P_{\{a, c\} \times \{b\}}$ and $\mathbb{P}(E_2) = P_{\{a, b\} \times \emptyset} - P_{\{a, b\} \times \{a, c\}}$. We now define a bijection between the events E_1 and E_2 by *exchanging the b content* of R_i and R_j , that is, by changing the sets R_i and R_j into R'_i and R'_j respectively, where $R'_i = R_i \cup \{b\}$ if $b \in R_j$ and $R'_i = R_i \setminus \{b\}$ otherwise, and similarly $R'_j = R_j \cup \{b\}$ if $b \in R_i$ and $R'_j = R_j \setminus \{b\}$ otherwise. Since exchanging the b content creates a bijection between the events E_1 and E_2 (which preserves the fact that (R_1, \dots, R_n) is in $\mathcal{M}_{p_1, p_2, p_3}^n$) the probabilities of the events E_1 and E_2 are equal. \square

Using Lemma 7.1 (three times) in (7.1) gives

$$\mathbb{P}(\text{tree}) = P_{\{1, 2\} \times \emptyset} + P_{\{1, 3\} \times \emptyset} + P_{\{2, 3\} \times \emptyset}.$$

Thus the probability that $G = \alpha(B)$ is a tree is equal to the probability that R_i has cardinality 2. Theorem 1.6 follows.




FIGURE 10. Events for which $G = \alpha(B)$ is a tree in the case $k = 3$ of Theorem 1.6. The notation $\bar{a} \in R$ means $a \notin R$.

Case $k = 4$ and greater. In the case $k = 4$ one considers $B = ((i, j, \ell), (R_1, \dots, R_n))$, where $i, j, \ell \in [n]$ and $(R_1, \dots, R_n) \in \mathcal{M}_{p_1, p_2, p_3, p_4}^n$. One gets the following analogue of (7.1):

$$\begin{aligned} \mathbb{P}(\text{tree}) = & \tilde{P}_{\{1\} \times \{2\} \times \{3\}} - \tilde{P}_{\{1\} \times \{2\} \times \{3,4\}} + \tilde{P}_{\{1\} \times \{2,3\} \times \{3,4\}} - \tilde{P}_{\{1\} \times \{2,3,4\} \times \{3\}} \\ & - \tilde{P}_{\{1,2\} \times \{2,3\} \times \{3,4\}} + \tilde{P}_{\{1\} \times \{4,1,2\} \times \{2,3\}} + \tilde{P}_{\{1\} \times \{4,1,2\} \times \{3,4\}} - \tilde{P}_{\{1\} \times \{4,1,2\} \times \{2,3,4\}} \\ & + \tilde{P}_{\{1,2,3\} \times \{3,4\} \times \{4,1\}} - \tilde{P}_{\{1,2\} \times \{2,3,4\} \times \{3,4,1\}} + \tilde{P}_{\{1,2,3\} \times \{3,4,1\} \times \{4,1,2\}} \end{aligned}$$

with $\tilde{P}_{A,B,C} = \sum_{t=1}^4 P_{t+A, t+B, t+C}$, where $t+A$ is the cyclically shifted set $\{t+a, a \in A\}$ and $P_{A,B,C}$ is the probability of the event $\{A \subseteq R_i, B \subseteq R_j \text{ and } C \subseteq R_\ell\}$. This expression can be obtained by direct inspection of the 16 possible Cayley trees, or by using the matrix-tree theorem. Using the later method we were able to compute analogues of (7.1) for all $k \leq 9$.

In the case $k = 4$ we can use *exchange lemmas* similar to Lemma 7.1 to prove Theorem 7.1 (one has to use six such lemmas, and the calculations can be organized according to the order of magnitude of these probabilities when n tends to infinity, with p_1, p_2, p_3, p_4 kept fixed). One can in fact generalize the type of exchange lemmas used, and use Gröbner bases to prove Theorem 1.6 (hence Theorem 1.2) up to $k = 7$. There is however one extra difficulty starting with $k = 4$: the *exchange operation* as the one used in the Proof of the lemma 7.1 become invalid unless one conditions on the fact that the indices i, j, ℓ, \dots under consideration are distinct (otherwise the exchange operation might affect one of the sets outside of the exchange). Thus for $k \geq 4$, the computer assisted proofs require to consider the different subcases corresponding to having non-distinct indices among i, j, ℓ, \dots . In a forthcoming paper [3] we shall give a general proof of Theorem 1.6 for all k and prove other similar results.

REFERENCES

- [1] O. Bernardi. An analogue of the Harer-Zagier formula for unicellular maps on general surfaces. *Adv. in Appl. Math.*, 48(1):164–180, 2012.
- [2] O. Bernardi and A.H. Morales. Counting cacti the easy way. In preparation.
- [3] O. Bernardi and A.H. Morales. Trees from sets: a probabilistic arborescence with algebraic roots. In preparation.
- [4] I. Goulden and A. Nica. A direct bijection for the Harer-Zagier formula. *J. Combin. Theory Ser. A*, 2005.
- [5] J. Harer and D. Zagier. The Euler characteristic of the moduli space of curves. *Invent. Math.*, 85(3):457–485, 1986.

- [6] D. M. Jackson. Some combinatorial problems associated with products of conjugacy classes of the symmetric group. *J. Combin. Theory Ser. A*, 49(2):363–369, 1988.
- [7] S.K. Lando and A.K. Zvonkin. *Graphs on surfaces and their applications*. Springer-Verlag, 2004.
- [8] B. Lass. Démonstration combinatoire de la formule de Harer-Zagier. *C. R. Acad. Sci. Paris*, 333:155–160, 2001.
- [9] B. Mohar and C. Thomassen. *Graphs on surfaces*. J. Hopkins Univ. Press, 2001.
- [10] A.H. Morales and E. Vassilieva. Bijective enumeration of bicolored maps of given vertex degree distribution. In *DMTCS Proceedings, 21st International Conference on Formal Power Series and Algebraic Combinatorics (FPSAC 2009)*.
- [11] A.H. Morales and E. Vassilieva. Direct bijective computation of the generating series for 2 and 3-connection coefficients of the symmetric group. *arXiv:1101.3614*, 2011.
- [12] G. Schaeffer and E. Vassilieva. Partitioned cacti: a bijective approach to the cycle factorization problem. In *DMTCS Proceedings, 20th Annual International Conference on Formal Power Series and Algebraic Combinatorics (FPSAC 2008)*.
- [13] G. Schaeffer and E. Vassilieva. A bijective proof of Jackson’s formula for the number of factorizations of a cycle. *J. Combin. Theory, Ser. A*, 115(6):903–924, 2008.
- [14] R.P. Stanley. *Enumerative combinatorics, volume 2*. Cambridge University Press, 1999.

Olivier Bernardi, Alejandro H. Morales, Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA USA 02139.

`{bernardi, ahmorales}@math.mit.edu`